STATUS OF HIGH PERFORMANCE COMPUTING WITHIN THE LRP

1 Summary of Original LRP Goals in Computation

Two *strong recommendations* were made as part of the Computation section:

1. the CADC host data archives of upcoming projects and work on new data-mining techniques;
2. funds be allocated toward the provision of a mid-range HPC facility to be hosted at CITA.

Two further *recommendations* were made:

1. NSERC equipment grants be increased to enable researchers to keep pace with the volume of data and computation;
2. a nationally-funded multidisciplinary HPC network be established through CFI funding.

The first *strong recommendation* is addressed in the report by David Schade on the CADC's activities and the status of data mining. In this report we will address the second *strong recommendation*, and discuss issues related to the two *recommendations.*

2 HPC in Contemporary Astronomy

2.1 Overview

Astronomy is in the midst of a silent revolution. Driven by the continued advancement of Moore's Law (that computational power and storage capacity double every two years) and improvements in detector technology, the size of observational data-sets is growing exponentially. To handle this flood of data, astronomers require simpler and faster data-access via international networks. When theory and simulation are included in this picture, it becomes clear that High Performance Computing (HPC) is an indispensable and enabling technology for modern astronomy and astrophysics.

Moreover, computation has become fully established as the third "tyne" of the astronomical research "trident" as illustrated in Figure 1. Interaction between theory and experiment often works indirectly through the computation "conduit". Furthermore, the development of the Virtual Observatory concept means that over the coming years, the relationship between computation and observation will become even closer. Advances in compute power, networking and visualization will only serve to cement this relationship.

2.2 Science Drivers

The four new major observatories planned for the next 20 years, ALMA, JWST, SKA and VLOT will provide data of unprecedented size and accuracy. While analysing the data from these instruments will be challenging in itself, the effect these instruments have on the development of theory will be profound. Astrophysical theory is currently undergoing a shift from single to multi-physics models, and the new data will serve to push this transformation. Science areas that will be heavily impacted by the new observatories, as well as advances in HPC, include:

- **Data analysis:** Using ALMA as an example, given a data-rate of 200 TB/year, the visibilities database is expected to exceed 1 PB of data after 5 years of observations.

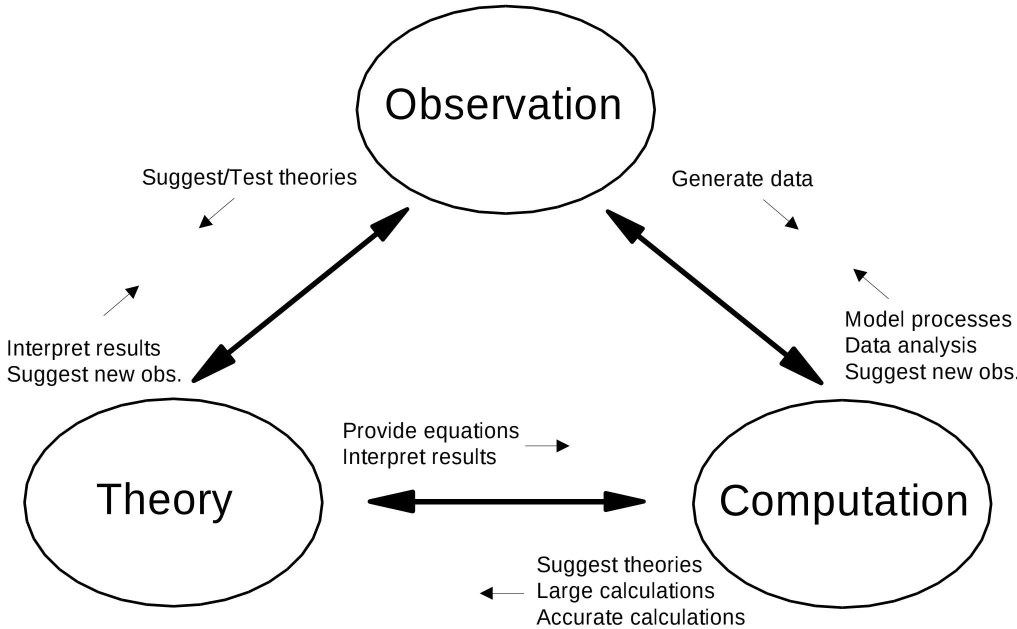

Figure 1: The “trident” of modern scientific investigation.

- **Planet Formation:** Data forthcoming from the ALMA project will place new constraints on theories of planet formation. A cohesive picture of planet formation needs to incorporate self-gravitational MHD, dust, and radiative transfer.

- **Star Formation:** While the qualitative picture of star formation has been known for decades, simulating the full collapse problem is exceptionally challenging. Not only is the radiation hydrodynamics problem difficult from a technical perspective (6-D), the change in length scales is extreme (seven orders of magnitude).

- **Supernovae:** At present the best ignition models are 2-D. Improving this to 3-D will require a 100-fold improvement in computational power.

- **Galaxy Formation and Evolution:** The combination of insufficient resolution, and a lack of a detailed understanding of global star formation make current models comparatively primitive.

- **Numerical Relativity:** The growth of gravitational wave astronomy will require precise calculations of wave signatures for various astrophysical phenomena. 3-D calculations of black hole coalescence are in their very earliest stages, and require a significant increase in compute power.

- **Astrophysical Turbulence:** Turbulence continues to be a topic that defies simple explanations. Astrophysical magnetic Reynolds numbers can exceed 10^{16}, whereas current simulations can reach effective Reynolds numbers of only a few thousand. Thus, there is still a great deal of resolution to be explored to cover realistically the cascade (and inverse cascade) of energy from large scales to small.

- **3-D Models of Stellar Structure:** A detailed understanding of convection zones is vital as they affect stellar structure, evolution, spectra and oscillations. A crucial step in stellar structure modelling is the addition of radiation hydrodynamic processes which are computationally expensive.

3 International Landscape

The top 500 supercomputers in the world are ranked according to computational power semi-annually (http://www.top500.org). From 2001 through to November 2003, Canada has had a maximum of five academic machines on this list, and a minimum of one. The list has a rapid turnover, both because of Moore's Law and a growth in the average number of CPUs per machine: machines are getting faster and *larger*. Matching an exponential profile to the growth of computing power shows that a machine placed at #1 initially, can be expected to take 7 years to fall off the list. Within the past few years, no academic machine in Canada has exceeded 30% of the power of the 10th placed machine on the list. In the last two years, Canada has ranked between 5th and 8th relative to the G8 countries.

Although no longer host to the largest machine in the world, the USA holds a sizable lead in terms of the number of supercomputers on the top 500. Even still, losing the #1 position on the top 500[1] has prompted a knee-jerk response in the US that is polarizing the American HPC community. The NSF Cyberinfrastructure report has recommended increased investment in HPC, although present budgetary constraints are severely hampering this push. Today, for example, creating a machine to occupy the #1 spot on the top 500 would cost in excess of 10^8 US$. However, a #10 machine would cost less than 10^7 US$. The message is clear: being the best is very expensive; being second best isn't.

Other countries are aggressively pursuing improved HPC infrastructure as well. Spain has recently invested US$80 million to build a new national HPC facility. In Germany, the Max Planck Institute for Plasma Physics (IPP) established a computing facility in 1960 which has evolved into the RZG Supercomputing Centre. The RZG is now so large it serves both the IPP and the Max Planck Institute for Astrophysics and regularly places in the top 20 of the top 500 (currently 31).

Numerous research groups around the world who do not have access to, or choose not to access nationally supported academic top 500 sites, have been able to access funds to purchase systems with several hundred CPUs for their own dedicated use. This is a concern because these efforts, which seem attractive to the funding agencies because of their relatively low costs, have the effect of concentrating HPC resources for use by the few. For example, in cosmology, a number of individual research groups now have dedicated access to large clusters (*e.g.*, Hernquist's group at Harvard, Frenk's group at Durham, Cambridge Relativity Group, and to a lesser extent the UK Astrophysical Fluids Facility at Leicester and even CITA's McKenzie). Research units with dedicated access to such large facilities have a unique advantage: they are able to run large scale simulations on these machines for weeks at a time. This situation has led to Hernquist's group taking a clear lead in simulations of cosmological structure formation by running a series of models which took 6 weeks each to run on their cluster.

4 National Issues

4.1 Funding

Since the LRP was written, the CFI and provincial funding agencies have committed almost $240 million to HPC. This investment has resulted in the establishment of 7 major regional computing collaborations in Canada: WestGrid (Alberta, B.C.), SHARCNET (Southern Ontario), PSciNet (Toronto), HPCVL (Eastern Ontario), RQCHP (Quebec), CLUMEQ (Eastern Quebec), and most recently, ACEnet (Atlantic Canada). As a result, many mid-range servers have been purchased for academic use across Canada, and the situation for computational science is significantly better in 2004 than it was in 1999. Through PSciNet, CITA purchased a 576-processor "Beowulf" cluster ("McKenzie"), as well as a 32 processor SMP. As one of four ACEnet sites, Saint Mary's University will install more than $5 million of new computational and visualization

[1] currently Japan's 'Earth Simulator', also dubbed 'Computenik' by US journalists in analogy with Sputnik.

equipment over the next two years to support, primarily, its new Institute for Computational Astrophysics (ICA). SHARCNET is currently preparing to place new servers at its member institutions (McMaster, UWO, Waterloo, Guelph, Windsor, Laurier, York, Brock, UOIT, Sheridan, Fanshawe) as part of a $49 million CFI2 grant. Further, since these installations are CFI-supported, at least 20% of the cycles on the new machines are available to outside users from any discipline (including astronomers at other institutions across Canada).

4.2 Installed Facilities

As of November 2003, 4 Canadian universities housed computing facilities that listed in the top 500. In order of ranking on the top 500 they are UBC (58), Toronto (CITA, 70; Particle Physics, 140), Calgary (450), and McGill (499). Facilities with a computational power in excess of 120 GFlops are present at Alberta, Queen's, Victoria, UWO and McMaster.

4.3 Canadian Strengths

Canadians have an excellent track record in computation. Significant contributions have been made to the development of cosmology, magneto-hydrodynamics, galaxy and planetary dynamics, and stellar structure and atmospheres. Canadian researchers have also contributed to the development of a number of key computational packages used by the astrophysics community.

4.4 Institute for Computational Astrophysics

Since the LRP was written in 1999, Saint Mary's University used some of its CRC complement to found the Institute for Computational Astrophysics (ICA). After an 18-month search, two Americans, as it turns out, were awarded CRCs in computational astrophysics: Dr. Robert Deupree from Los Alamos, well known for his work in the Helium Flash problem and multi-dimensional stellar structure, accepted a tier I chair and assumed the ICA directorship, and Dr. Joe Hahn from the Lunar and Planetary Institute in Houston accepted a tier II chair. Deupree and Hahn joined two existing computational faculty at SMU (Clarke and Guenther), as well as a new faculty appointment created by the university and accepted by Dr. Ian Short (part of the Phoenix team on stellar atmospheres) who is a Maritimer by birth, and who was in a tenure-track position in Florida. Since then, two post-doctoral fellows have been appointed to the ICA as part of a five-year plan to acquire a steady-state PDF population of eight, and a new Ph.D. programme in Astronomy has been approved and implemented at Saint Mary's, the first science Ph.D. programme in Nova Scotia outside Dalhousie.

Computational resources at the fledgling ICA were modest during the time the three new faculty were being recruited, and the importance of its participation in ACEnet in attracting the new faculty and PDFs to Halifax cannot be overstated. While much of this report appropriately emphasizes the need for a nationally funded, top 20 academic site, the roles played by well-supported regional sites such as ACEnet should not be forgotten. Indeed, national sites could well be built on the shoulders of vigorous regional sites who already have the building and staff support structure in place.

The LRP report in 1999, which was enthusiastically endorsed by Saint Mary's president, played a strong role in exciting the senior administration about establishing the ICA, and it is fair to say that this report is in part responsible for the institutes's existence. In as much as the ICA then helped make ACEnet a more credible proposal, the LRP can take some indirect credit for the success of ACEnet. However, ACEnet itself, which had very strong participation from the physics, chemistry, computing science, and oceanography communities in Atlantic Canada, would have been funded without the LRP (though SMU's and thus astronomy's participation would have been significantly reduced).

4.5 CITA

CITA was a member of PSciNet I, and used these funds to first acquire a 32 processor alpha GS320 shared memory machine, as well as a 16 processor alpha cluster. This investment was supported by the university through the provision of a parallel programmer.

While this initial infrastructure was expensive per unit of raw performance, it allowed researchers to rapidly exploit the availability of a significant number of compute cycles with minimal changes to existing codes. Scientific results came out rapidly. As experience with parallel systems grew, some production codes were ported to run effectively on cheaper distributed memory commodity hardware platforms. When the time came to upgrade the infrastructure, a decision was made to transition to a massively parallel commodity processor cluster (McKenzie) which ranked 34th on the top 500 list in June 2003. By using commodity parts, the system was assembled at a modest cost of under $1 million (CAD). A number of codes were able to successfully exploit the new hardware within weeks of arrival, and translate the world class horsepower into leading edge scientific results.

Throughout this period, CITA has also supported the Canadian astronomy community. As part of a commitment to CFI, 20% of cycles on CITA facilities are held for external users. Massively parallel applications currently running on McKenzie include; SPH for cosmology and planet formation, MHD for black hole accretion, supernovae, cosmology and interstellar medium simulations, linear algebra for CMB analysis, Monte-Carlo applications ranging from stellar structure to chemical engineering.

A key theme at CITA has been the focus on leading edge computation, and a significant fraction of the cycles are used in dedicated mode (‘capability computing’) where a single application can obtain months of the whole machine to perform the worlds biggest simulations. The CITA model has also shown that a science-focused machine can be installed and maintained at very low overhead.

5 Summary

Neither the ICA nor the regional HPC consortia such as ACEnet, SHARCNET, and WestGrid, who directly and indirectly support computational astrophysicists, were part of the landscape when the 1999 LRP report was released. Thus, the strong recommendation to support a mid-range system at CITA now seems out of date and under-ambitious. Nonetheless, through a proposal that originated from the University of Toronto (not a national effort), CITA is now host to the McKenzie cluster as well as a 32-processor SMP. While to some extent this can be viewed as having realized the strong recommendation, external access is limited to only 20% of the research time on these machines, and, more importantly, these machines were designed and installed primarily with the needs of CITA researchers in mind, not the national community.

The CFI fund-and-match formula has resulted in the establishment of regional consortia, where mid-range computing is now readily available. Conversely, national facilities have not emerged due to the provincial-match funding mechanism. Regional facilities, that are shared among many disciplines and researchers, whilst providing resources for most mid-scale users who need dozens of processors, are often heavily utilized. This places users who wish to use the entire resource for a particular project at a clear disadvantage; these systems are rarely free for such endeavours. With the establishment of ACURA it is now possible to present a national vision for top tier infrastructure. A cooperative national effort could procure a machine with thousands of processors that focuses on running large jobs in a dedicated mode, enabling science that cannot be realized on any mid-range facility.

6 Future Proposals

Astronomers collaborate well. The necessity of cooperative proposals for off-shore installations has developed a culture with very little geographic tension, and a great motivation to succeed. In stark contrast, the preparation of an LRP for HPC in Canada, led by C3.ca, has proven to be exceedingly challenging, primarily because of the inability of different fields to see beyond their own individual needs and aspirations. Further, it has become clear that astronomers in Canada have an exceptional HPC skill set that appears to lead other disciplines.

In light of recent developments (the ICA and regional consortia), funding a nationally-supported machine, perpetually updated to remain in the top 20, is realistic. Housing such a facility at one of the regional consortia, where building infrastructure and support staff are already in place, would be prudent economically. The astronomical community may wish to suggest that such a machine be acquired for the exclusive use by astronomers, but a coordinated effort with the greater HPC community in Canada (*e.g.*, the LRP for HPC effort) is clearly an alternative approach. However, in any cooperative effort time-sharing issues will be central to the success of the proposal.

Recommendation: A top 20 facility be installed and maintained in perpetuity. Such a facility could be installed at an individual institution, at one or more of the regional HPC consortia, or possibly at a specifically constructed national site in collaboration with interested scientific communities (*e.g.*, materials physics, chemistry, and biology). The inclusion of Boards of Governors of HPC consortia, university authorities, federal funding agencies (NRC, NSERC, CFI or its successor), and possibly ACURA, will be a necessary part of the proposal process. Collaboration with C3.ca, and inclusion in the LRP for HPC, is desirable but not necessary. Costs can be expected to be in the region of $15 million (CAD) per 3 year technology refresh cycle.

The irrepressible advance of the top 500 means that astronomers in Canada need uninhibited access to an internationally competitive facility within the next three years to prevent our recent strides forward from being eroded.

Robert Thacker, Ue-Li Pen, and David Clarke